

## Claims

1           1.    A method of topology propagation in a distributed  
2   computing environment, said method comprising:

3                sending group connectivity messages from at least  
4   one group leader to identified nodes of at least one  
5   group of nodes within the distributed computing  
6   environment;

7                discontinuing said sending of group connectivity  
8   messages during a time period of no topology change  
9   within the distributed computing environment; and

10               reinitiating sending of group connectivity  
11   messages from the at least one group leader upon  
12   identification of a topology change within the  
13   distributed computing environment.

1           2.    The method of claim 1, wherein the distributed  
2   computing environment comprises at least two networks each  
3   having at least one group of identified nodes, and wherein  
4   said method further comprises employing within each group of  
5   the at least two networks a heartbeat protocol to ensure  
6   continued presence of each identified node within the group.

1           3.    The method of claim 2, wherein the at least two  
2   networks of the distributed computing environment comprise  
3   heterogenous networks.

1           4.    The method of claim 2, wherein at least one node  
2   of the distributed computing environment has at least two  
3   adapters, said at least two adapters coupling said at least  
4   one node to said at least two networks, and wherein said  
5   sending comprises sending first group connectivity messages  
6   (GCMs) from a first group leader to identified nodes of a  
7   first group of nodes on a first network of said at least two  
8   networks, said at least one node comprising an identified  
9   node of said first group of nodes, and forwarding said first  
10  GCMs by said at least one node to a second group of nodes on  
11  a second network of said at least two networks.

1           5.    The method of claim 4, wherein said first GCMs  
2   received at identified nodes of said first group of nodes  
3   and identified nodes of said second group of nodes are  
4   employed by each said identified node to update a local  
5   network connectivity table (NCT).

1           6.    The method of claim 4, wherein said sending  
2   further comprises sending second GCMs from a second group  
3   leader to identified nodes of the second group of nodes, and  
4   forwarding said second GCMs by said at least one node to the  
5   first group of nodes on the first network of the at least  
6   two networks.

1           7.    The method of claim 6, wherein said sending second  
2   GCMs by said second group leader is responsive to receiving  
3   new information in said forwarded first GCMs at said second  
4   group leader.

1           8.    The method of claim 6, wherein said discontinuing  
2 comprises for each group leader discontinuing said sending  
3 of group connectivity messages when a number of messages  
4 sent from the group leader reaches a set limit after  
5 identification by said group leader of a topology change  
6 within the distributed computing environment.

1           9.    The method of claim 8, wherein said reinitiating  
2 comprises identifying said topology change within a  
3 distributed computing environment, said identifying  
4 comprising at least one of: receiving at a group leader a  
5 node connectivity message which conflicts with a local  
6 network connectivity table value, receiving at a group  
7 leader a group connectivity message which conflicts with a  
8 local network connectivity table value, identifying that a  
9 local adapter belongs to a different adapter membership  
10 group, or identifying that a local adapter has become  
11 disabled.

1           10.   The method of claim 1, wherein said discontinuing  
2 comprises for each group leader discontinuing said sending  
3 of group connectivity messages when a number of messages  
4 sent from the group leader reaches a set limit after  
5 identification of the topology change within the distributed  
6 computing environment.

1           11.   The method of claim 1, further comprising  
2 implementing said sending, said discontinuing, and said  
3 reinitiating without employing acknowledgment messages  
4 during said topology propagation.



1           13. A system for topology propagation in a distributed  
2 computing environment, said system comprising:

3           means for sending group connectivity messages from  
4 at least one group leader to identified nodes of at  
5 least one group of nodes within the distributed  
6 computing environment;

7           means for discontinuing said sending of group  
8 connectivity messages during a time period of no  
9 topology change within the distributed computing  
10 environment; and

11           means for reinitiating sending of group  
12 connectivity messages from the at least one group  
13 leader upon identification of a topology change within  
14 the distributed computing environment.

1           14. The system of claim 13, wherein the distributed  
2 computing environment comprises at least two networks each  
3 having at least one group of identified nodes, and wherein  
4 said system further comprises means for employing within  
5 each group of the at least two networks a heartbeat protocol  
6 to ensure continued presence of each identified node within  
7 the group.

1           15. The system of claim 14, wherein the at least two  
2 networks of the distributed computing environment comprise  
3 heterogenous networks.

1        16. The system of claim 14, wherein at least one node  
2 of the distributed computing environment has at least two  
3 adapters, said at least two adapters coupling said at least  
4 one node to said at least two networks, and wherein said  
5 means for sending comprises means for sending first group  
6 connectivity messages (GCMs) from a first group leader to  
7 identified nodes of a first group of nodes on a first  
8 network of said at least two networks, said at least one  
9 node comprising an identified node of said first group of  
10 nodes, and means for forwarding said first GCMs by said at  
11 least one node to a second group of nodes on a second  
12 network of said at least two networks.

1        17. The system of claim 16, wherein said first GCMs  
2 received at identified nodes of said first group of nodes  
3 and identified nodes of said second group of nodes are  
4 employed by each said identified node to update a local  
5 network connectivity table (NCT).

1        18. The system of claim 16, wherein said means for  
2 sending further comprises means for sending second GCMs from  
3 a second group leader to identified nodes of the second  
4 group of nodes, and means for forwarding said second GCMs by  
5 said at least one node to the first group of nodes on the  
6 first network of the at least two networks.

1        19. The system of claim 18, wherein said means for  
2 sending second GCMs by said second group leader is  
3 responsive to receiving new information in said forwarded  
4 first GCMs at said second group leader.

1        20. The system of claim 18, wherein said means for  
2        discontinuing comprises for each group leader means for  
3        discontinuing said sending of group connectivity messages  
4        when a number of messages sent from the group leader reaches  
5        a set limit after identification by said group leader of a  
6        topology change within the distributed computing  
7        environment.

1        21. The system of claim 20, wherein said means for  
2        reinitiating comprises means for identifying said topology  
3        change within a distributed computing environment, said  
4        means for identifying being responsive to at least one of:  
5        receiving at a group leader a node connectivity message  
6        which conflicts with a local network connectivity table  
7        value, receiving at a group leader a group connectivity  
8        message which conflicts with a local network connectivity  
9        table value, identifying that a local adapter belongs to a  
10       different adapter membership group, or identifying that a  
11       local adapter has become disabled.

1        22. The system of claim 13, wherein said means for  
2        discontinuing comprises for each group leader means for  
3        discontinuing said sending of group connectivity messages  
4        when a number of messages sent from the group leader reaches  
5        a set limit after identification of the topology change  
6        within the distributed computing environment.

1        23. The system of claim 13, wherein said means for  
2        sending, said means for discontinuing, and said means for  
3        reinitiating are implemented without employing  
4        acknowledgment messages during said topology propagation.

1        24. The system of claim 13, wherein said means for  
2        reinitiating sending of group connectivity messages is  
3        responsive to at least one of receiving at a group leader a  
4        node connectivity message which conflicts with a local  
5        network connectivity table value, receiving at a group  
6        leader a group connectivity message which conflicts with a  
7        local network connectivity table value, identifying that a  
8        local adapter belongs to a different adapter membership  
9        group, or identifying that a local adapter has become  
10       disabled.



1        25. At least one program storage device readable by a  
2 machine tangibly embodying at least one program of  
3 instructions executable by the machine to perform a method  
4 of topology propagation in a distributed computing  
5 environment, comprising:

6            sending group connectivity messages from at least  
7 one group leader to identified nodes of at least one  
8 group of nodes within the distributed computing  
9 environment;

10           discontinuing said sending of group connectivity  
11 messages during a time period of no topology change  
12 within the distributed computing environment; and

13           reinitiating sending of group connectivity  
14 messages from the at least one group leader upon  
15 identification of a topology change within the  
16 distributed computing environment.

1        26. The at least one program storage device of claim  
2 25, wherein the distributed computing environment comprises  
3 at least two networks each having at least one group of  
4 identified nodes, and wherein said method further comprises  
5 employing within each group of the at least two networks a  
6 heartbeat protocol to ensure continued presence of each  
7 identified node within the group.

1        27. The at least one program storage device of claim  
2 26, wherein the at least two networks of the distributed  
3 computing environment comprise heterogenous networks.

1        28. The at least one program storage device of claim  
2 26, wherein at least one node of the distributed computing  
3 environment has at least two adapters, said at least two  
4 adapters coupling said at least one node to said at least  
5 two networks, and wherein said sending comprises sending  
6 first group connectivity messages (GCMs) from a first group  
7 leader to identified nodes of a first group of nodes on a  
8 first network of said at least two networks, said at least  
9 one node comprising an identified node of said first group  
10 of nodes, and forwarding said first GCMs by said at least  
11 one node to a second group of nodes on a second network of  
12 said at least two networks.

1        29. The at least one program storage device of claim  
2 28, wherein said first GCMs received at identified nodes of  
3 said first group of nodes and identified nodes of said  
4 second group of nodes are employed by each said identified  
5 node to update a local network connectivity table (NCT).

1        30. The at least one program storage device of claim  
2 28, wherein said sending further comprises sending second  
3 GCMs from a second group leader to identified nodes of the  
4 second group of nodes, and forwarding said second GCMs by  
5 said at least one node to the first group of nodes on the  
6 first network of the at least two networks.

1        31. The at least one program storage device of claim  
2 30, wherein said sending second GCMs by said second group  
3 leader is responsive to receiving new information in said  
4 forwarded first GCMs at said second group leader.

1        32. The at least one program storage device of claim  
2 30, wherein said discontinuing comprises for each group  
3 leader discontinuing said sending of group connectivity  
4 messages when a number of messages sent from the group  
5 leader reaches a set limit after identification by said  
6 group leader of a topology change within the distributed  
7 computing environment.

1        33. The at least one program storage device of claim  
2 32, wherein said reinitiating comprises identifying said  
3 topology change within a distributed computing environment,  
4 said identifying comprising at least one of: receiving at a  
5 group leader a node connectivity message which conflicts  
6 with a local network connectivity table value, receiving at  
7 a group leader a group connectivity message which conflicts  
8 with a local network connectivity table value, identifying  
9 that a local adapter belongs to a different adapter  
10 membership group, or identifying that a local adapter has  
11 become disabled.

1        34. The at least one program storage device of claim  
2 25, wherein said discontinuing comprises for each group  
3 leader discontinuing said sending of group connectivity  
4 messages when a number of messages sent from the group  
5 leader reaches a set limit after identification of the  
6 topology change within the distributed computing  
7 environment.

1        35. The at least one program storage device of claim  
2 25, further comprising implementing said sending, said  
3 discontinuing, and said reinitiating without employing  
4 acknowledgment messages during said topology propagation.

1        36. The at least one program storage device of claim  
2 25, wherein said reinitiating sending of group connectivity  
3 messages comprises at least one of receiving at a group  
4 leader a node connectivity message which conflicts with a  
5 local network connectivity table value, receiving at a group  
6 leader a group connectivity message which conflicts with a  
7 local network connectivity table value, identifying that a  
8 local adapter belongs to a different adapter membership  
9 group, or identifying that a local adapter has become  
10 disabled.

\* \* \* \* \*